# PRELIMINARY

# Product Requirement Document

## Reference Design for
## Wired μSAN Storage Appliance

Revision 0.8

Marketing____________________Date_________

Engineering____________________Date_________

Arcitecture____________________Date_________

Software____________________Date_________

CTO____________________Date_________

no 13
3603388

## Contents

# 1 Introduction

This document provides the requirements for a μSAN Storage Appliance (hereafter referred to as simply μSAN). The intent of this document is to specify those requirements for operation, performance, environment and maintenance that are believed to be key for successful development and implementation of the product prototype.

# 2 Product Definition

The μSAN is an Internetworking Protocol (IP) aware storage appliance that uses a revolutionary new communications protocol developed by Zetera Corporation. The μSAN Storage Appliance provides inexpensive, scalable storage to any network utilizing IP communications protocol regardless of the medium used to connect network elements. The key to μSAN being so versatile is that it uses standard IP communications protocol that is supported by virtually every PC, switch and network router. In addition, μSAN supports an ever increasing number of IP aware Consumer Electronics devices are coming on the market.

The major function of the μSAN is to allow many connected peers to fairly share the storage resources contained within. Each peer can own a portion of the μSAN storage, called a partition, as if that storage existed within the peer itself. Each peer communicates with its partition by directly addressing it with a unique IP address. Ownership of storage partitioned in this way effectively amortizes the cost of the μSAN among the connected peers. In addition, the μSAN provides for strong security, reliability, and scalability. Any μSAN can be spanned to other μSANs to increase the size of a partition. Several μSANs can be used as mirrors or RAID. μSAN arrays provide performance, redundancy and massive amounts of capacity.

The μSAN defined by this document consists of a high performance, high capacity 7200 RPM EIDE/ATA hard disk drive with an IDE to μSAN protocol adapter (Tailgate) enclosed in a plastic enclosure powered by an external power supply. It is very similar in appearance and cost to existing 1394/USB external drives, but is vastly more capable.

# 3 Market Overview

Consumer Electronics (CE) has long been confronted with the need to address mass storage due to the transport of large amounts of video and audio data. We see this in the desire for PVR (personal video recorder) and MP3 jukebox functionality. Consumers expect high functionality, reliability and low cost in high volume CE products. However, at the same time, CE manufacturers and suppliers cannot overcome basic cost obstacles to achieve high unit volume. We also see the need for additional storage and storage access for many IP aware CE applications. Clearly the CE market is moving towards the use of massive numbers of Hard Disk Drives (HDDs).

Even though the cost per gigabyte on HDDs decreases significantly every year, as areal bit density increases, the lowest unit cost of a HDD is virtually fixed. The current trend in areal density is driving the size of even the smallest HDDs to levels far larger than many small CE devices could use. A consumer electronics device does not need a $100, 100 gigabyte solution so much as it needs a $10, 10-gigabyte solution or a $1, 1-gigabyte solution.

An obvious solution to this cost/size problem is to disaggregate the drive, create separate SKUs, and to amortize the cost of the storage across many low-cost device applications. This is precisely what a μSAN does. A single μSAN can service the needs of many consumer electronics devices. The μSAN protocol allows the requesting device to be as simple or complex as the application needs further lowering the cost. The μSAN is a low-cost product. As such, it provides a compelling cost advantage to CE devices that need to share data.

# 4 Key Features

The following is a list of key features of the μSAN storage appliance:

- IP connected storage appliance
- Directly supported by 802.3 compliant devices
- Ethernet 10/100 mb/s RJ-45 connector for network attachment
- Block address processing performed at full wire speed
- Fully NAT compatible (routable)
- Storage capacities from 80GB to 200GB
- Storage element is a 7200 RPM EIDE HDD
- Independent partitions sharable by connected peers
- Auto initialization from first power-up
- Supports Universal Plug and Play
- Automatic volume spanning (virtually unlimited drive size)
- Automatic mirroring support
- Fanless for quiet operation
- B/W allocation and other QoS features supported
- OS independent – works with Windows 98/ME/NT/2000/XP, Linux, Mac O/S, Unix, Solaris, and many other operating systems
- Supports multiple levels of security
- Supports Unicast and Multicast communications
- Root level setup is fully autonomous – requires no master controller
- Supports volume sharing among connected peers
- Supports persistent volume attributes – access levels and authorities
- Supports temporary volume attributes and ownership
- Supports Block I/O data transfers with low overhead
- Provides all advantages of DAS on a network

# 5 Financial Considerations

The cost of storage, although widely considered to be a bargain at today's capacities, is still a major component of the total cost of computing devices and storage nodes in a networked system. As the home becomes more and more connected using IP aware devices, storage costs will be a major consideration there as well. Therefore, it is absolutely imperative that the cost of ownership of a μSAN storage appliance be as low as possible.

The following are μSAN storage appliance cost targets:

| | |
|---|---|
| μSANTailgate PCB | $15.00 |
| Power Supply Brick (12VDC) | $5.00 |
| Plastic Enclosure | $5.00 |
| Retail box, manual, CD, etc | $2.50 |
| Assembly and test | $2.50 |
| Shipping and Handling | $3.00 |
| **Total μSAN COG$ [1]** | **$33.00** |

A typical ASP$ for a wired 80GB μSAN would be:

| | | |
|---|---|---|
| 80 GB HDD [2] | $60.00 | |
| μSAN COG$ | $33.00 | |
| OEM Margin$ (20-30%) | $23.25 | $ 39.86 |
| OEM Price to Retailer | $116.25 | $ 132.86 |
| Retail Margin$ (30-40%) | $49.82 | $88.57 |
| Retail ASP$ | $166.07 | $221.43 |
| **Retail Shelf Price** | **$169.95** | **$224.95** |

Note [1]: This figure includes a negotiated μSAN license fee.

Note [2]: This price assumes a high volume purchase agreement with HDD OEM.

# 6 Functional Specifications

## 6.1 System level

| | |
|---|---|
| Communications Standards | IEEE 802.3, 802.3u |
| Protocol | CSMA/CD |
| LAN port | RJ- 45 connector |
| Speed | 10/100(half duplex), 20/200(full duplex) |
| Indicators | Power/Activity/status Red / Yellow / Green LED |
| Controls | Power On/Off rocker switch on rear panel |
| Reset Switch | Recessed PB switch on rear panel |
| Cabling | UTP Category 5 minimum |

## 6.2 Hard Disk Storage

| | |
|---|---|
| Storage type | EIDE/ATA Hard Disk Drive |
| HDD Capacity[1] | 80GB, 120, 180, 200 GB |
| HDD Rotational speed | 7200 RPM |
| Ave HDD latency[1] | 4.2 msec |
| Ave HDD seek time[1] | 8.5 msec |
| Native data rate | UDMA 66/100 |
| Operational protocol | µSAN V1.0 |
| Power Supply | 12V DC external modular supply |
| Power connector | 3.0 mm barrel connector |

Note [1]: These values may vary according to manufacturer's specifications

## 6.3 Control electronics

| | |
|---|---|
| Control PCB | 4-layer PCB attached to HDD |
| Processor | ARM-7 embedded processor |
| Onboard memory | 16 MB flash ROM |
| | 32 MB DRAM |
| PCB I/O | 40-pin EIDE UDMA 66/100 |
| | Standard RJ-45 802.3/3u |
| Power connector | 3.0mm female barrel connector |
| Onboard P/S | 5V +/- 5% @ 2A HDD logic supply |
| | 3.3V +/- 5% @ 2A logic supply |

## 6.4 External power source

| | |
|---|---|
| Power supply (standard) | External wall-plug with 2 m pwr cable |
| Input power | 110 VAC @ 60Hz (standard) |
| Power supply (world wide) | External brick, adapter cable and 2 m pwr cable |
| Input power | 90-240 VAC @ 50-60Hz |
| Output power | +12V +/- 10% @ 3A - continuous |
| | +12V+/-20% @ 5A – 10 sec max (spin up) |

# 7 Performance Specifications

## 7.1 Transfer Command Overhead

512 byte transfer efficiency of (512/588) = 87%

ACK CMD Overhead

Total Number of bytes = 72

## 7.2 Request Command Overhead

Total Number of bytes = 76

Ethernet Inter-frame Gap

Ethernet inter-frame gap 12 bytes.

## 7.3 Write Data Transfer Performance *

| | |
|---|---|
| 1Mbyte from Requester to µSAN: = | 136.8 ms @ 100Mb/s<br>13.6 ms @ 1Gbe |
| Maximum throughput = | 7.31 MBytes/Sec on 100Mb/s<br>73.1 MBytes/Sec on 1Gbe |

## 7.4 Read Data Transfer Performance*

| | |
|---|---|
| 1MByte from the µSAN = | 137.6 ms @ 100Mb/s<br>13.6 MS @ 1Gbe |
| Maximum throughput = | 7.27 MBytes/Sec<br>72.7 MBytes/Sec) |

*** assumes no latency, no collision**

# 8 Environmental Specifications

| | |
|---|---|
| Operating temp | 0°C to 40°C (32°F to 104°F) |
| Storage temp | -20°C to 70°C (-4°F to 158°F) |
| Temperature gradient | <3°C/minute |
| Operating humidity | 10% to 85% Non-condensing |
| Storage humidity | 5% to 90% Non-condensing |
| Operating altitude | -100M to 3000M |
| Operating shock | <20G 2ms ½ sine wave - any axis |
| Non-op shock | <200G 2 ms ½ sine wave - any axis |

# 9 Physical Specifications

| | |
|---|---|
| Dimensions | Height 1.74 in. (44.20mm) nom |
| | Length 8.61 in. (218.75mm) nom |
| | Width 6.08 in. (154.43mm) nom |
| Weight (without P/S) | 1.82 lb (827 gm) |

# 10 Agency and Certification Requirements

| | |
|---|---|
| Safety Certification | UL 1950 |
| | CSA-C22.2 No 950 |
| Electromagnetic | FCC Class B |
| | CSA Class B |

# 11 Software Functional Requirements

## 11.1 O/S support

### 11.1.1 Linux Red Hat V 8.0 running on embedded ARM7 H/W platform

### 11.1.2 Wind River RTOS

### 11.1.3 PSOS

## 11.2 Support d Commands

### 11.2.1 Transfer Data

When issued this command the μSAN will transfer the data either as a write to the μSAN or the result of a request from the μSAN. One block of data is transferred. In a write from Requester operation, this is the only command that is transferred from the Requester. The μSAN responds with an ACK Command.

### 11.2.2 Request Data

When issued this command, the μSAN responds with a Transfer Command containing the requested data. No ACK Command is required in response to this transfer.

### 11.2.3 Request Lock Command

This command is a request for a transfer from the μSAN, however, it locks out all other accesses from any other authenticated Requester. The lock is released when a Transfer Command is received from this Requester or a time-out occurs. The timer for the time-out is retriggerable upon Request Command accesses from the initiating Requester. Time-out TBD. The function of this command is to provide for semaphore functionality of the data.

### 11.2.4 ACK Command

This command acknowledges a successful transfer. It need only be used in when the communication is being executed through the UDP transport layer to recognize a successful Transfer command from a Requester to the μSAN. Even in a TCP environment, the Requester needs to await an ACK before proceeding to insure single threaded error exception handling. The LBA that was received in the Transfer Command is sent to resolve the "Sorcerer's Apprentice" problem inherent in IP due to time-out retry.

### 11.2.5 Error Command

The Error Command indicates that an operation could not be completed for some reason and is a transfer from the μSAN to the Requester.

Error Codes:

| | |
|---|---|
| 01 | Invalid Authorization |
| 02 | Partition has locked you out |
| 04 | Go Command had Invalid Authorization at destination |
| 08 | Go Command locked out of Partition at the destination |
| 10 | LBA out of Range |
| 20 | LBA is Write Protected |

### 11.2.6 Go Transfer Command

The Go Transfer Command informs a µSAN to go active on the network and autonomously transfer a block of data to another µSAN. The purpose of this command is to execute a copy to another µSAN without going through the Requester. The result is to double the effective bandwidth and to remove a burden from the Requester. This does not turn the µSAN into a master, the Name to IP resolution is still performed by the Requester and an ACK is communicated back to the Requester upon completion. The µSAN uses TCP to transport this data to the second µSAN.

### 11.2.7 Go Request Command

The Go Request Command is similar to the Go Transfer Command. The Requester issues this command to a µSAN to have it request data from another µSAN. At the completion of the transfer, the µSAN will ACK or ERROR the Requester. TCP transport is used to send the request to the second µSAN.

### 11.2.8 Go Request Lock Command

The Go Request Lock Command is similar to the Go Request Command with the lock element added to it.

### 11.2.9 Release Partition Command

This command will cause the µSAN to release the partition, erase the data blocks, and release the IP/Name placing the blocks back into the root pool for future allocation.